

NASA/CR-2001-210870
ICASE Report No. 2001-12




An Assessment of Linear Versus Non-linear Multigrid Methods for Unstructured Mesh Solvers

Dimitri J. Mavriplis
ICASE, Hampton, Virginia

ICASE
NASA Langley Research Center
Hampton, Virginia

Operated by Universities Space Research Association



National Aeronautics and
Space Administration

Langley Research Center
Hampton, Virginia 23681-2199

Prepared for Langley Research Center
under Contract NAS1-97046

May 2001

Form SF298 Citation Data

Report Date <i>("DD MON YYYY")</i> 00MAY2001	Report Type N/A	Dates Covered (from... to) <i>("DD MON YYYY")</i>
Title and Subtitle An Assessment of Linear Versus Non-linear Multigrid Methods for Unstructured Mesh Solvers		Contract or Grant Number
		Program Element Number
Authors Dimitri J. Mavriplis		Project Number
		Task Number
		Work Unit Number
Performing Organization Name(s) and Address(es) Aeronautics and Space Administration Langley Research Center Hampton, Virginia 23681-2199		Performing Organization Number(s)
Sponsoring/Monitoring Agency Name(s) and Address(es)		Monitoring Agency Acronym
		Monitoring Agency Report Number(s)
Distribution/Availability Statement Approved for public release, distribution unlimited		
Supplementary Notes ICASE Report No. 2001-12		
Abstract Abstract. The relative performance of a non-linear FAS multigrid algorithm and an equivalent linear multigrid algorithm for solving two different non-linear problems is investigated. The first case consists of a transient radiation-diffusion problem for which an exact linearization is available, while the second problem involves the solution of the steady-state Navier-Stokes equations, where a first-order discrete Jacobian is employed as an approximation to the Jacobian of a second-order accurate discretization. When an exact linearization is employed, the linear and non-linear multigrid methods converge at identical rates, asymptotically, and the linear method is found to be more efficient due to its lower cost per cycle. When an approximate linearization is employed, as in the Navier-Stokes cases, the relative efficiency of the linear approach versus the non-linear approach depends both on the degree to which the linear system approximates the full Jacobian as well as the relative cost of linear versus non-linear multigrid cycles. For cases where convergence is limited by a poor Jacobian approximation, substantial speedup can be obtained using either multigrid method as a preconditioner to a Newton-Krylov		
Subject Terms		

Document Classification unclassified	Classification of SF298 unclassified
Classification of Abstract unclassified	Limitation of Abstract unlimited
Number of Pages 22	

AN ASSESSMENT OF LINEAR VERSUS NON-LINEAR MULTIGRID METHODS FOR UNSTRUCTURED MESH SOLVERS

DIMITRI J. MAVRIPLIS*

Abstract. The relative performance of a non-linear FAS multigrid algorithm and an equivalent linear multigrid algorithm for solving two different non-linear problems is investigated. The first case consists of a transient radiation-diffusion problem for which an exact linearization is available, while the second problem involves the solution of the steady-state Navier-Stokes equations, where a first-order discrete Jacobian is employed as an approximation to the Jacobian of a second-order accurate discretization. When an exact linearization is employed, the linear and non-linear multigrid methods converge at identical rates, asymptotically, and the linear method is found to be more efficient due to its lower cost per cycle. When an approximate linearization is employed, as in the Navier-Stokes cases, the relative efficiency of the linear approach versus the non-linear approach depends both on the degree to which the linear system approximates the full Jacobian as well as the relative cost of linear versus non-linear multigrid cycles. For cases where convergence is limited by a poor Jacobian approximation, substantial speedup can be obtained using either multigrid method as a preconditioner to a Newton-Krylov method.

Key words. unstructured, multigrid, Krylov

Subject classification. Applied and Numerical Mathematics

1. Introduction. Multigrid methods are well known as efficient solution techniques for both linear and non-linear problems. As with many iterative solvers, multigrid methods can be used directly as non-linear solvers [1, 4, 5], or as linear solvers operating on a linearization arising from a Newton solution strategy for the non-linear problem at hand [25, 3, 18]. In addition, multigrid can also be used as a linear or non-linear preconditioner for a Newton-Krylov method [14, 2, 15].

Newton solution strategies for non-linear problems incorporating linear multigrid solvers may fail when the initial guess is far removed from the domain of convergence of the non-linear problem, and globalization methods may be required to ensure a convergent method. Non-linear multigrid methods overcome this difficulty by using a pseudo-time-stepping analogy on the non-linear problem directly [1, 5]. On the other hand, non-linear multigrid methods may fail due to the non-existence of a solution to the physical problem which is rediscretized on the coarse grid levels, particularly in the initial stages of convergence. However, for various applications such as time-dependent problems, where the initial guess provided from the previous time step is often within the non-linear convergence domain of the next time step, or steady-state problems with mild non-linearities such as subsonic or transonic flows (as opposed to hypersonics), these issues are often of minor importance.

Non-linear multigrid methods require the evaluation of the full non-linear residual at each iteration on all grid levels, while linear multigrid methods replace these operations by matrix (Jacobian) vector products at each iteration on all grid levels, with the evaluation of non-linear residuals only occurring on the fine grid

*ICASE, Mail Stop 132C, NASA Langley Research Center, Hampton, VA 23681-2199, U.S.A., dimitri@icase.edu. This research was partially supported by the National Aeronautics and Space Administration under NASA Contract No. NAS1-97046 while the author was in residence at ICASE, NASA Langley Research Center, Hampton, VA 23681-2199. Partial support was also provided under U.S. Department of Energy subcontract B347882 from Lawrence Livermore National Laboratory.

at each outer Newton iteration. One of the great advantages of non-linear multigrid methods is that they obviate the need to form and store the Jacobian matrix associated with the Newton strategy. For many large-scale unstructured mesh computations, where memory is the limiting factor, non-linear multigrid methods are indeed the only viable solution strategies [13]. On the other hand, in cases where the non-linear residual evaluation is costly, linear multigrid methods may become more attractive on a cpu-time efficiency basis, since for a fixed stencil, the cost of the Jacobian-vector products is fixed and independent of the cost of non-linear residual evaluations, the latter of which are only performed a small number of times in the outer Newton iteration. Of course this statement is only true provided the convergence of both methods is similar on a multigrid iteration basis. In the asymptotic convergence region, where solution updates become small, and the effect of non-linearities vanishes, it can be shown, and has been observed, that both methods converge at the same rates per multigrid cycle, provided equivalent iteration strategies are used in both cases (linear and non-linear Jacobi for example).

The above discussion is only valid in the case where an exact Newton linearization of the non-linear problem is employed in the linear multigrid method, and an exact local linearization is used in the non-linear method. For discretizations which are not confined to nearest-neighbor stencils, such as second-order accurate convection operators which rely on distance-two neighbor stencils, the evaluation and storage costs of the exact Jacobian become prohibitive, and simpler Jacobians based on first-order accurate nearest neighbor stencils are most often employed. This practice, which can be thought of as a defect-correction scheme or a preconditioning approach [11], ensures that quadratic convergence of the outer Newton iteration will never be achieved, and hence that solution of the linear system to high tolerances even in the asymptotic convergence range will be fruitless. Therefore, the overall solution efficiency of a non-linear multigrid method versus a linear multigrid method in such cases depends not only on the relative cost of non-linear residual evaluations versus Jacobian-matrix vector products, but also on the degree to which a partial solution of the reduced Jacobian system is successful in converging the full non-linear system.

In the following paper we examine two problems which are solved with a non-linear full approximation storage (FAS) multigrid method [1], a linear multigrid method, and multigrid preconditioned Newton-Krylov methods. The first problem is a transient two-equation radiation diffusion model which contains strong non-linearities, but for which an exact Jacobian can easily be constructed. The second problem is the solution of the steady-state Euler and Navier-Stokes equations. In this case, the non-linearities are less pronounced for the flow regimes considered than in the radiation problem, but a first-order accurate Jacobian is used to solve the second-order accurate discretization, for the reasons described above. While these two test problems serve to demonstrate two different situations for the comparison of linear versus non-linear multigrid methods, the eventual solution of coupled radiation-hydrodynamic systems is also of interest.

2. Linear and Non-Linear MG Formulations. The goal of any multigrid method is to accelerate the solution of a fine grid problem by computing corrections on a coarser grid and then interpolating them back to the fine grid problem. Although this procedure is described in a two grid context, it is applied recursively on a complete sequence of fine and coarser grid levels. To apply a linear multigrid method to a non-linear problem, a linearization must first be performed. Thus, if the equations to be solved are written as

$$(2.1) \quad \mathbf{R}_h(\mathbf{w}_{\text{exact}}) = \mathbf{0}$$

with the current estimate \mathbf{w} yielding the non-linear residual \mathbf{r} :

$$(2.2) \quad \mathbf{R}_h(\mathbf{w}_h) = \mathbf{r}$$

the Newton linearization of this system is taken as

$$(2.3) \quad \frac{\partial \mathbf{R}_h}{\partial \mathbf{w}_h} \Delta \mathbf{w}_h = -\mathbf{r}$$

This represents a linear set of equations in the solution variable $\Delta \mathbf{w}_h$ (the correction), to which a linear multigrid (i.e. MG correction scheme) can be applied. In this case, the coarse grid equation reads:

$$(2.4) \quad \frac{\partial \mathbf{R}_H}{\partial \mathbf{w}_H} \Delta \mathbf{w}_H = -I_h^H \mathbf{r}_{\text{linear}}$$

where H and h represent coarse grid and fine grid values, respectively, and I_h^H represents the restriction operator which interpolates the fine grid residuals to the coarse grid. The residual of the linear system on the fine grid is given by

$$(2.5) \quad \mathbf{r}_{\text{linear}} = \frac{\partial \mathbf{R}_h}{\partial \mathbf{w}_h} \Delta \mathbf{w}_h + \mathbf{r}$$

and may be approximated as

$$(2.6) \quad \mathbf{r}_{\text{linear}} \approx \mathbf{R}(\mathbf{w} + \Delta \mathbf{w})$$

where \mathbf{R}_h refers to the non-linear residual, as previously. The coarse grid corrections $\Delta \mathbf{w}_H$ which are obtained by solving equation (2.4) are initialized on the coarse grid as zero. After the solution of equation (2.4), these corrections are prolonged or interpolated back to the fine grid.

Alternatively, a non-linear FAS multigrid scheme can be used to solve equation (2.1) directly without resorting to a linearization. In this case, the FAS coarse grid equation reads:

$$(2.7) \quad \mathbf{R}_H(w_H) = \mathbf{R}_H(\tilde{I}_h^H w_h) - I_h^H \mathbf{r}$$

where the term on the right-hand side is often referred to as the defect-correction [1, 11]. \mathbf{R}_H represents the coarse grid discretization and I_h^H and \tilde{I}_h^H denote the restriction operators which are now used to interpolate residuals as well as flow variables from the fine grid to the coarse grids. In principal, different restriction operators for residuals and variables may be employed. If equation (2.7) is re-written as:

$$(2.8) \quad \mathbf{R}_H(w_H) - \mathbf{R}_H(\tilde{I}_h^H w_h) = -I_h^H \mathbf{r}$$

the right hand sides of equations (2.4) and (2.8) represent similar approximations of the restricted non-linear residual, in view of equation (2.6) and the fact that these restricted residuals in the FAS scheme are always evaluated at the most recently available fine grid updates. Therefore, by equating the left hand sides of equations (2.4) and (2.8), the equivalence between the linear multigrid scheme and the non-linear FAS scheme is seen to be given by:

$$(2.9) \quad \mathbf{R}_H(w_H) - \mathbf{R}_H(\tilde{I}_h^H w_h) \approx \frac{\partial \mathbf{R}_H}{\partial \mathbf{w}_H} \Delta \mathbf{w}_H$$

which means that the FAS multigrid scheme corresponds to an approximation to a linear multigrid scheme, where the coarse grid Jacobians are approximated by finite differencing the operator. Therefore, in the limit of asymptotic convergence, i.e. when $\Delta \mathbf{w}_H \ll 1$, the two methods should yield similar convergence rates.

Note that the above discussion involves no specification of the coarse grid operator and Jacobian construction. Therefore, a fair comparison of linear versus non-linear multigrid methods should utilize a similar construction for both of these quantities in the respective algorithms.

3. Multigrid Algorithms. The two multigrid variants implemented in this work are based on the agglomeration multigrid strategy. Agglomeration multigrid was originally developed for finite-volume schemes [7, 22, 30], and is based on agglomerating or fusing together neighboring fine grid control-volumes to form larger coarse grid control volumes as depicted in Figure 3.1. This approach has since been generalized for arbitrary discretizations following algebraic multigrid principles [9]. In fact, agglomeration multigrid can be viewed as a simplification and extension of algebraic multigrid to non-linear systems of equations. The control-volume agglomeration algorithm can be recast as a graph algorithm, similar to algebraic multigrid methods, where the “seed” vertex initiating an agglomerated cell corresponds to a coarse grid point, and the neighboring agglomerated points correspond to fine grid points, in the algebraic multigrid terminology [20]. While weighted graph algorithms can be employed for agglomeration, these weights cannot depend on solution values, as in the algebraic multigrid case, but only on grid metrics. In this manner, the coarse grid levels are static and need only be constructed at the beginning of the simulation.




FIG. 3.1. *Illustration of Agglomeration Multigrid Coarse Level Construction*

As in the algebraic multigrid case, agglomeration multigrid employs a Galerkin projection for the construction of the coarse grid equations. Thus, the coarse grid operator is given by:

$$(3.1) \quad \mathbf{R}_H = I_h^H \mathbf{R}_h I_H^h$$

where I_h^H is the restriction operator, and I_H^h is the prolongation operator, and both operators are taken as piecewise constants. This simple construction applies equally to linear and non-linear operators, and reduces to forming the coarse grid equation at an agglomerated cell as the sum of the fine grid equations at each fine

grid cell contained in the coarse grid cell. The non-linearities in the operator are evaluated using solution variables on the coarse grid interpolated up from the fine grid.

Given this multigrid infrastructure, two particular algorithms which differ mainly in the manner in which non-linearities are handled are developed for comparison. The first involves a standard non-linear FAS multigrid algorithm, and the second involves a linear multigrid algorithm applied to the linearization of the governing equations.

3.1. FAS Scheme. In the non-linear FAS multigrid algorithm, equation (2.1) is solved directly. The coarse grid equations are formed by Galerkin projection (c.f. equation (3.1)) and the non-linearities in the coarse grid operator are evaluated using coarse level solution variables interpolated up from the fine grid using the \tilde{I}_h^H restriction operator (as per equation (2.8)). On each grid level, the discrete equations are solved using a Jacobi preconditioned multi-stage time-stepping scheme (for the Navier-Stokes equations) [16, 27, 26, 10] or a non-linear block Jacobi iteration which can be written as:

$$(3.2) \quad \mathbf{w}^{\text{new}} = \mathbf{w}^{\text{old}} + [D]^{-1} \mathbf{R}(\mathbf{w}^{\text{old}})$$

where $[D]$ represents the block diagonal of the Jacobian matrix. This smoother constitutes a non-linear solver, since the non-linear residual is updated at each stage, and incurs minimum memory overheads since only the storage of the block matrix $[D]$ representing the coupling between the solution variables at each grid point is required. This scheme is equivalent to a single stage Jacobi preconditioned multi-stage time-stepping scheme.

3.2. Linear Multigrid Scheme. The linear multigrid scheme solves equation (2.3) on the fine grid, and equation (2.4) on the coarse levels. On the fine grid, the Jacobian $\frac{\partial \mathbf{R}_h}{\partial \mathbf{w}_h}$ is formed by explicitly differentiating (hand coding) the discrete operator \mathbf{R}_h . On the coarse levels, for consistency with the FAS multigrid algorithm, the Jacobian is taken as the explicit differentiation of the coarse non-linear operator obtained by Galerkin approximation (c.f. equation (3.1)). Thus flow variables as well as residuals are restricted to the coarser grids, but the non-linear residuals on these coarser levels are not evaluated, only the Jacobians corresponding to the linearization of the non-linear coarse level residuals. These coarse level Jacobians are evaluated at the beginning of the solution phase for the non-linear time-step problem, and are then held fixed throughout the linear multigrid iterations. The multi-level linear system constructed in this manner more closely approximates the equivalent FAS scheme, as opposed to the more traditional approach of agglomerating the fine grid Jacobian terms directly. Memory requirements for the linear multigrid scheme are increased over those of the FAS scheme due to the required storage of the fine and coarse level Jacobians.

An outer Newton iteration is employed to solve the complete non-linear problem $\mathbf{R}(w) = 0$. Within each Newton iteration, the linear system defined by equation (2.3) is solved by the linear multigrid algorithm. This provides a fine grid correction Δw which is then used to update the non-linear residual. These non-linear iterations converge quadratically provided the linear system is solved to sufficient tolerance and a consistent linearization is employed. When approximate Jacobian representations are employed, such as in the Navier-Stokes equations, slower convergence of this outer iterative procedure is obtained.

On each grid level, the linear multigrid scheme solves the linear system using a block-Jacobi smoother. If the Jacobian is divided up into diagonal and off-diagonal block components, labeled as $[D]$ and $[O]$, respectively, the Jacobi iteration can be written as:

$$(3.3) \quad [D] \Delta \mathbf{w}_h^{n+1} = -\mathbf{r} - [O] \Delta \mathbf{w}_h^n$$

where $\Delta \mathbf{w}_h^n$ represents corrections from the previous linear iteration, and $\Delta \mathbf{w}_h^{n+1}$ represents the new linear corrections produced by the current linear iteration. At each linear iteration, the solution of equation (3.3) requires the inversion of the block matrix $[D]$ at each grid point. The linear corrections $\Delta \mathbf{w}_h$ are initialized to zero at the first iteration on each grid level. Therefore, this linear iteration strategy reduces to the non-linear Jacobi scheme described above in the event only a single linear iteration is employed.

In contrast to the non-linear FAS multigrid algorithm, the residuals, jacobians (i.e. $[D]$ and $[O]$ terms), and the variables interpolated up to the coarse grids are only evaluated at the start of the non-linear iteration, and are held fixed during all inner linear multigrid cycles within a non-linear iteration.

4. Radiation Diffusion Problem. The non-equilibrium radiation diffusion equations can be written as

$$\frac{\partial E}{\partial t} - \nabla \cdot (D_r \nabla E) = \sigma_a (T^4 - E) \quad (4.1)$$

$$\frac{\partial T}{\partial t} - \nabla \cdot (D_t \nabla T) = -\sigma_a (T^4 - E)$$

with

$$\sigma_a = \frac{z^3}{T^3}, \quad D_r(T, E) = \frac{1}{3\sigma_a + \frac{1}{E} \left| \frac{\partial E}{\partial n} \right|}, \quad D_t(T) = \kappa T^{\frac{5}{2}}$$

Here, E represents the photon energy, T is the material temperature, and κ is the material conductivity. In the non-equilibrium case, the non-linear source terms on the right-hand-side are non-zero and govern the transfer of energy between the radiation field and material temperature. Additional non-linearities are generated by the particular form of the diffusion coefficients, which are functions of the E and T variables. In particular, the energy diffusion coefficient, $D_r(T, E)$ contains the term $\left| \frac{\partial E}{\partial n} \right|$ which refers to the gradient of E in the direction normal to the cell interface (in the direction of the flux). This limiter term is an artificial means of ensuring physically meaningful energy propagation speeds (i.e. no larger than the speed of light) [2, 6, 14]. The atomic number z is a material coefficient, and while it may be highly variable, it is only a function of position (i.e. $z = f(x, y)$ in two dimensions).

Equations (4.1) represent a system of coupled non-linear partial-differential equations which must be discretized in space and time. Spatial discretization on two-dimensional triangular meshes is achieved by a Galerkin finite-element procedure, assuming linear variations of E and T over a triangular element. The non-linear diffusion coefficients are evaluated by first computing an average T and E value along a triangle edge, and then computing the non-linear diffusion coefficient at the edge midpoint using these averaged values. The gradient of E in the D_r diffusion coefficient is also taken as a one dimensional gradient along the direction of the stencil edge. The source terms are evaluated using the local vertex values of E and T exclusively, rather than considering linear variations of these variables.

The time derivatives are discretized as first-order backwards differences, with lumping of the mass matrix, leading to an implicit scheme which requires the solution of a non-linear problem at each time step. This approach is first-order accurate in time, and is chosen merely for convenience, since the principal objective is the study of the solution of the non-linear system.

The Jacobian of the required linearizations is obtained by differentiation (hand coding) of the discrete non-linear residual. Because the spatial discretization involves a nearest neighbor stencil, the Jacobian can be expressed on the same graph as the residual discretization, which corresponds to the edges of the triangular grid. The initial guess for the solution of the non-linear problem at each time-step is taken as the solution obtained at the previous time-step.

The test case chosen for this work is taken from [14] and depicted in Figure 4.1. We consider a unit square domain of two dissimilar materials, where the outer region contains an atomic number of $z = 1$ and the inner regions $(1/3 < x < 2/3), (1/3 < y < 2/3)$ contains an atomic number of $z = 10$. The top and bottom walls are insulated, and the inlet and outlet boundaries are specified using mixed (Robin) boundary conditions, as shown in the figure. This domain is discretized using a triangular grid containing 7,502 vertices, shown in Figure 4.2. This grid conforms to the material interface boundaries in such a way that no triangle edges cross this boundary.

Figure 4.3 illustrates a typical simulation for this case. Incoming radiation sets up a traveling thermal front in the material, the progress of which is impeded by the region of higher atomic number z . At critical times in the simulation, the diffusion coefficients can vary by up to six orders of magnitude near the material interfaces, thus providing a challenging non-linear behavior for the multigrid algorithms. At each physical time step, a non-linear problem must be solved. It is the solution of this transient non-linear problem at a given time step which forms the test problem for the two agglomeration multigrid algorithms. Clearly, the size of the physical time step affects the stiffness of the non-linear problem to be solved, with smaller physical time-steps leading to more rapidly converging systems. The non-dimensional time-step chosen in this simulation was taken as 0.01. This constitutes a rather large value compared to those employed in reference [14] (usually of the order of 10^{-3}) and may have an adverse effect on overall temporal accuracy, but provides a more stringent test case for the multigrid solvers. Of the order of 1000 time steps are required to propagate the thermal front from the inlet to outlet boundary in the current simulation.




FIG. 4.1. *Sample test problem for non-linear radiation-diffusion equations*




FIG. 4.2. *Illustration of unstructured grid for non-linear radiation-diffusion problem: 7,502 vertices*




FIG. 4.3. *Illustration of solution for non-linear radiation-diffusion problem: Contours of T*

Table 4.1 depicts the relative time required for a non-linear residual evaluation on the fine grid, assembly of the various Jacobian matrix entries on the fine grid, and timings for various components of the linear and non-linear multigrid algorithms. The residual and Jacobian terms are assembled within the same loop for cache efficiency reasons, and minimal incremental work is incurred for computing the additional off-diagonal Jacobian terms, required for the linear multigrid scheme. This is due to the fact that much of the block diagonal (point) Jacobian terms consist of the sum of the corresponding off-diagonal Jacobian terms, and thus require the same computations. For both linear and non-linear schemes, the block diagonal Jacobians must be inverted, as shown in equations (3.2) and (3.3). For multiple Jacobi sweeps, the LU decomposition of these block matrices is formed on the first pass, and then frozen for subsequent passes. Thus the first linear or non-linear Jacobi iteration incurs additional cost over subsequent passes, as depicted in the table. The timings illustrate the lower cost of the linear iterations, which are up to five times faster than the corresponding non-linear iterations. In the non-linear case, the initial iteration involves the computation of a non-linear residual, the diagonal Jacobian terms, and the LU decomposition of these Jacobians, while subsequent iterations only require the evaluation of the non-linear residuals. In the linear case, the first iteration includes the LU decomposition of the point Jacobians, but does not include residual and Jacobian construction timings, (which are relegated to the outer Newton iteration). From the table, the non-linear FAS multigrid cycle is seen to require four times more cpu time than the equivalent linear multigrid cycle. In this case, a 4-level W(3,0) saw-tooth cycle was used, with three (linear or non-linear) Jacobi iterations performed on each level when going from fine to coarse levels. These timings do not include the outer Newton iteration in the linear case, which incurs a non-linear residual evaluation and Jacobian construction, noting that this expense may be amortized over a variable number of linear multigrid cycles.

TABLE 4.1
*Relative CPU Time Required for Various Components
of Linear and Non-Linear Multigrid Methods for Radiation
Diffusion Problem*

Component	Normalized Timing
Non-Linear Residual	1.0
Residual + Point Jacobians	2.52
Residual + Entire Jacobian	2.82
1st Stage Non-Linear Sweep	2.82
Add. Stages Non-Lin. Sweeps	1.07
1st Linear Jacobi Sweep	0.364
Incr. Linear Jacobi Sweeps	0.173
FAS MG Cycle	13.04
Linear MG Cycle	3.31




FIG. 4.4. *Convergence Rate for Transient Radiation Problem in terms of Multigrid Cycles (3 linear MG cycles per Newton update)*




FIG. 4.5. *Convergence Rate for Transient Radiation Problem in terms of Multigrid Cycles (5 linear MG cycles per Newton update)*




FIG. 4.6. *Convergence Rate for Transient Radiation
Problem in terms of Normalized CPU Time*

Figure 4.4 provides a comparison of the convergence rates of the linear and non-linear multigrid schemes in terms of the number of multigrid cycles. The four-level W(3,0) saw-tooth cycle described above is employed in both cases. For the linear scheme, three linear multigrid cycles are employed for each Newton update, and the convergence history of the linear residual is plotted alongside that of the non-linear residual. As expected, quadratic convergence of the non-linear residual is initially observed in the linear-multigrid-Newton scheme. However, this quadratic behavior is lost after four Newton iterations, due to the inexact solution of the linear problem (using three multigrid cycles). In this region, the convergence rate of the outer non-linear Newton scheme becomes governed or limited by the convergence of the inner linear problem. In the quadratic convergence region, each time the non-linear residual is updated, the linear residual increases slightly, before resuming its downwards trend. In the asymptotic region, the linear and non-linear residuals become approximately equal, as expected from equation (2.6). In this region, the convergence of the non-linear FAS multigrid scheme and the linear multigrid scheme become equivalent, in terms of the number of multigrid cycles, as expected from equation (2.9). The fact that the linear multigrid convergence plot lies slightly to the right of the FAS convergence curve is due to the additional effort spent solving the linear system in the initial phases of slow non-linear convergence. Figure 4.5 further illustrates this point, by comparing the convergence history for the same approach using five linear multigrid cycles per Newton iteration. In this case, the final asymptotic convergence rate is similar, but is reached at a later stage and with additional numbers of multigrid sweeps due to increased oversolving of the linear system in the initial stages of non-linear convergence. Adaptive convergence criteria for the linear system can clearly aid in reducing oversolution of the linear system, although this has not been considered in this work.

Figure 4.6 compares the convergence efficiencies of the non-linear FAS multigrid approach with the linear multigrid approach using three multigrid cycles per Newton iteration, in terms of cpu time. The linear multigrid method is over three times more efficient, which can be entirely attributed to the lower cost per multigrid cycle of the linear multigrid scheme.

5. Solution of Steady-State Navier-Stokes Equations. The Navier-Stokes equations are discretized on mixed triangular-quadrilateral meshes using a vertex-based approach where the flow variables are stored at the grid vertices. Median-dual control volumes are constructed around each vertex, and fluxes at control volume interfaces are evaluated using a Roe approximate Rieman solver [19]. Second-order accuracy

is obtained through a simplified gradient reconstruction technique which results in a distance-two neighbor stencil. Viscous terms are constructed as diffusion operators involving nearest neighbor stencils. For inviscid flow simulations, the viscous fluxes are neglected, while for viscous turbulent flows these terms are retained, and the influence of turbulence is simulated using the Spalart-Allmaras one equation turbulence model [23]. The turbulence equation is discretized in the same manner as the flow equations, with the exception that the convective terms are only first-order accurate. The turbulence equation is solved simultaneously but uncoupled from the flow equations using the same multigrid algorithm.

The non-linear FAS multigrid solver employs a multi-stage time-stepping scheme as a smoother on all grid levels, which requires the evaluation of the non-linear residual at each stage. While the fine grid equations are discretized to second-order accuracy, the coarse level equations are only discretized to first-order accuracy. This simplifies their implementation on the coarse level agglomerated graphs, and is consistent with practices used on structured geometric multigrid solvers for similar problems [5, 9]. Local preconditioning is applied to the multistage scheme by pre-multiplying the non-linear residual by the inverted block diagonal Jacobian matrix at each stage [16, 27, 26, 10]. This is equivalent to a (scaled) non-linear Jacobi iteration at each stage. For viscous flows, line preconditioning is employed, which involves inverting the block tridiagonal Jacobian entries along lines constructed in boundary layer regions (c.f. Figure 5.5) [10, 12]. This corresponds to a non-linear iterative line solution technique which can be described by equation (3.2) where $[D]$ now represents the line Jacobians, instead of the diagonal elements. In isotropic grid regions, the lines reduce to a single point and the line preconditioning becomes equivalent to Jacobi preconditioning. In all cases, the local Jacobian entries correspond to those derived from a first-order discretization. The LU decomposition of these local Jacobians is performed on the first stage of the multi-stage scheme and then frozen for the subsequent stages of the scheme, thus amortizing the LU decomposition cost over multiple stages.

The linear multigrid method operates on the discrete Jacobian of the first-order discretization of the non-linear flow equations, although the fine grid flow equations are discretized to second-order accuracy. Non-linear residuals are only evaluated on the fine grid at the beginning of each linear solution phase, which may involve multiple linear multigrid sweeps. Coarse grid Jacobians are obtained by linearizing the non-linear coarse grid agglomerated operator, in order to provide a more consistent comparison between equivalent linear and non-linear methods. On each grid level, multiple passes of a linear Jacobi or Gauss-Seidel smoother are employed for inviscid flows. For viscous flows, multiple passes of a linear line solver are employed, following equation (3.3), where $[D]$ corresponds to the block tridiagonal Jacobians taken along the set of lines constructed in the grid, and $[O]$ corresponds to the remaining Jacobian entries. The Jacobi implementations of point and line algorithms correspond to the linear counterparts of the non-linear smoothers used in the FAS multigrid algorithm. In the Gauss-Seidel implementation, lines and points are pre-sorted in increasing x-direction, and sweeps using latest available updates are performed on the grid in increasing and decreasing x direction at odd and even smoothing passes, respectively. When multiple linear smoothing passes are employed, the LU decomposition of the local point or line Jacobians is performed on the first pass and then frozen for subsequent passes.

TABLE 5.1
*Relative CPU Time Required for Various Components
of Linear and Non-Linear Multigrid Methods for Inviscid
Fluid Flow Problem*

Component (Euler)	Normalized Timing
Non-Linear Residual	1.0
Residual + Line Jacobians	1.62
Residual + Entire Jacobian	1.86
1st Stage Non-Linear Sweep	1.96
Add. Stages Non-Lin. Sweeps	1.26
1st Linear GS Sweep	0.43
Incr. Linear GS Sweeps	0.38
3-stage FAS MG Cycle	8.92
5-stage FAS MG Cycle	9.86
Linear MG Cycle (1 W cycle)	5.70
Linear MG Cycle (2 W cycles)	8.98

TABLE 5.2
*Relative CPU Time Required for Various Components
of Linear and Non-Linear Multigrid Methods for Viscous
Fluid Flow Problem*

Component (Navier-Stokes)	Normalized Timing
Non-Linear Residual	1.17
Residual + Line Jacobians	2.13
Residual + Entire Jacobian	2.39
1st Stage Non-Linear Sweep	2.44
Add. Stages Non-Lin. Sweeps	1.44
1st Linear GS Sweep	0.43
Add. Linear GS Sweeps	0.38
3-stage FAS MG Cycle	10.4
5-stage FAS MG Cycle	11.3
Linear MG Cycle (1 W cycle)	6.3
Linear MG Cycle (2 W cycles)	9.6

For the inviscid case, an isotropic coarsening strategy which results in a coarsening ratio of 4:1 is employed for generating coarse agglomerated levels, while a directional coarsening strategy which proceeds in the direction of the implicit lines is employed in the viscous flow cases, also yielding a 4:1 reduction in complexity between fine and coarse levels [10, 12].

Tables 5.1 and 5.2 depict the relative cpu times required for the various components of the linear and non-linear algorithms on the grid of Figure 5.4, for both inviscid Euler computations and viscous Navier Stokes computations. Both a three stage [10, 28] and a five stage [8] non-linear smoother are examined. The time required for the 5-stage smoother increases only moderately over that required for the 3-stage scheme, due to the fact that the local Jacobian LU decomposition is only performed once for each scheme, and the dissipative terms are only evaluated three times for both schemes (at odd stages only for the 5-stage scheme) [8]. Assembly of the complete Jacobian required for the linear scheme incurs relatively little extra overhead over that required for assembling the point or line Jacobians, since many of the same terms required for the point Jacobians can be used in these additional off-diagonal Jacobian elements. Both linear and non-linear smoothers involve additional startup cost on the first smoothing pass, due to the need to perform the LU decomposition of the local Jacobians, which are frozen on subsequent passes. The Jacobi (not shown) and Gauss-Seidel variants of each linear solver are approximately equivalent in overall cost, and can be seen to be approximately four times less costly than the equivalent non-linear solver, mainly due to the fact that these smoothers avoid evaluation of the non-linear residual. Overall, a non-linear update using a single linear multigrid W-saw-tooth-cycle, with 4 Gauss-Seidel smoothing sweeps on each grid level, requires approximately 60% of the effort of a three-stage FAS scheme. Using two linear multigrid cycles per non-linear update results in a non-linear update cost approximately equal to that observed with the 3-stage FAS scheme for the inviscid flow case. The linear method efficiency advantage is slightly higher in the viscous flow case, since the non-linear residual is now augmented by the additional viscous terms which must be computed, while the linear smoother remains identical in cost since the stencil is unchanged from the inviscid case.

The inviscid test case consists of flow over a NACA 0012 airfoil at a Mach number of 0.8 and an incidence of 1.25 degrees. This well known test case produces a strong upper surface shock and a weaker lower surface shock. The unstructured triangular mesh for this case contains a total of 7,884 vertices. Figure 5.1 depicts the observed convergence rates for this case with the various schemes discussed above, compared in terms of cpu time required for a given level of reduction in the rms average of density residuals. A four level W-cycle was used for both multigrid schemes. The figure shows the linear multigrid approach, using a single W-cycle with four Gauss-Seidel smoothing passes, is approximately twice as efficient as the three-stage FAS scheme. The increase in efficiency between the Gauss-Seidel and Jacobi linear multigrid algorithms is due to the superior convergence properties of Gauss-Seidel over Jacobi, since both sweeps require approximately the same amount of cpu time. The 5-stage FAS scheme is slightly more efficient than the 3 stage scheme, as much of the local Jacobian LU decomposition and multigrid overhead is amortized over more grid sweeps.

Figure 5.2 illustrates the non-linear convergence rates achieved for the linear multigrid scheme per non-linear update, as a function of the number of linear multigrid cycles. The non-linear convergence rate has a lower bound which is approached as the number of linear multigrid cycles is increased and the linear system is solved more exactly. This asymptotic rate, which is in the neighborhood of 0.78, can be viewed as a measure to which the first-order Jacobian approximates the second-order discretization. (Quadratic convergence would be observed for an exact match). Note that only two linear W-cycles are effective at achieving most of this non-linear convergence, although the scheme using a single linear multigrid cycle is the most efficient overall, as shown in Figure 5.1. Improving the convergence past this threshold cannot be achieved with better linear solvers, but only through a more accurate Jacobian representation.

One way to achieve this is to use a matrix-free Newton-Krylov method [21, 15, 29] to approximate the exact Jacobian of the full second-order accurate residual. The linear multigrid solver provides a natural

candidate for a preconditioner of the Newton-Krylov method. The non-linear GMRES routine developed by Wigton and Yu [31] is employed for this purpose. This approach, which corresponds to a left-preconditioning strategy [21], also allows the use of a non-linear solver as a preconditioner, and hence the FAS multigrid solver is also implemented as a preconditioner for GMRES. Very rapid convergence in terms of non-linear updates is observed in Figure 5.2 when the Newton-Krylov method (using ten search directions) is applied with the linear multigrid method as a preconditioner. Figure 5.3 illustrates the convergence obtained with both multigrid schemes employed as solvers and as preconditioners for GMRES, using ten search directions, in terms of cpu time. The improvement is less dramatic when measured in this manner, since each non-linear update involves ten multigrid cycles. However, the Newton-Krylov method provides similar overall gains in efficiency for both the linear and non-linear schemes, particularly in the asymptotic convergence region.




FIG. 5.1. *Convergence Efficiency for Various Multigrid Algorithms for Inviscid Flow Problem*




FIG. 5.2. *Non-Linear Convergence Rate for Various Levels of Linear System Solution and for Linear MG Newton-Krylov Method*




FIG. 5.3. *Acceleration Provided by Newton-Krylov Method for Inviscid Flow Problem*

The next test case involves the computation of viscous turbulent transonic flow over an RAE 2822 airfoil at a Mach number of 0.73, and incidence of 2.31 degrees, and a Reynolds number of 6.5 million on the grid depicted in Figure 5.4. This grid contains a total of 16,167 vertices, and makes use of quadrilaterals in the highly stretched boundary layer and wake regions, and triangles in isotropic regions. The linear and non-

linear line algorithms are used in this case on the set of lines depicted in Figure 5.5, which were constructed using a previously developed graph algorithm [10]. The flowfield was initialized with freestream conditions and the turbulence model is converged simultaneously with the flow equations. Figure 5.6 illustrates the overall convergence of the various algorithms versus the number of non-linear iterations. In this case, a single linear W-cycle using four Jacobi smoothing passes on each level provides an asymptotically faster convergence rate per cycle than either FAS scheme, while the Gauss-Seidel version of this scheme is even faster. When these schemes are compared in terms of cpu time in Figure 5.7, the linear Gauss-Seidel scheme is three times more efficient than either non-linear FAS scheme, due to the superior convergence rate, as well as the lower cost per cycle of the linear multigrid scheme.




FIG. 5.4. *Illustration of Unstructured Grid for Viscous Flow over Airfoil (16,167 points)*




FIG. 5.5. *Illustration of Line Structure for Line Solver for Viscous Flow over Airfoil*




FIG. 5.6. *Convergence Efficiency of Various Algorithms in Terms of Outer Iteration Cycles for Viscous Airfoil Flow Problem*




FIG. 5.7. *Convergence Efficiency for Various Algorithms in Terms of CPU time for Viscous Airfoil Flow Problem*




FIG. 5.8. *Efficiency Gains of Newton-Krylov Method for Linear and Non-Linear Multigrid Algorithms for Viscous Airfoil Flow Problem*

Figure 5.8 illustrates the increased convergence efficiency when using the linear or non-linear multigrid scheme as a preconditioner for GMRES, employing ten search directions. Similar increases in convergence efficiency are obtained in both cases, with the non-linear scheme benefiting slightly more than the linear scheme. However, the linear multigrid preconditioned GMRES approach remains the overall most efficient solution technique.

The final test case involves subsonic viscous flow over a multi-element airfoil. The grid and associated line system are depicted in Figures 5.9 and 5.10. This mesh contains a total of 61,104 vertices, with quadrilateral elements in the boundary layer and wake regions, and triangular elements elsewhere. The Mach number for this case is 0.2, the incidence is 16 degrees, and the Reynolds number is 9 million. For this case, the Rieman solver is modified according to the low-Mach number preconditioning techniques developed previously [10, 26, 17]. The final computed solution in terms of Mach number contours is depicted in Figure 5.11. Complex cases of this nature have proved to be the most difficult to converge efficiently in past studies [10]. A 5 level W-cycle is used in all cases for the multigrid algorithms. The flowfield is initialized with a pre-converged solution obtained after 150 cycles of the FAS multigrid scheme (itself initialized from freestream conditions), and the turbulence model is frozen at its final converged values throughout these computations. This is done in order to focus on the asymptotic convergence behavior of the linear versus non-linear methods, and to avoid the complications of non-linear continuation which are required in this case for the linear solver operating on a freestream initialization.

In Figure 5.12 the relative convergence efficiencies of the linear and non-linear methods are displayed, as a function of the number of non-linear iterations. Although the linear multigrid scheme using a single W-cycle (with 4 Gauss-Seidel smoothing passes) initially converges faster than the 3-stage FAS scheme, the latter achieves a slightly faster asymptotic rate of convergence. However, both methods are almost equivalent asymptotically in terms of cpu time, since the linear multigrid method provides lower cost multigrid sweeps, as shown in Figure 5.13. In both cases, the overall convergence rate is over six times slower than that observed in the previous two cases. Solving the linear system to completion at each non-linear update (using 20 linear W-cycles) produces no observable benefit in non-linear convergence rate, as depicted in Figure 5.12. This indicates that the slower convergence in this case is attributable to a poor approximation of the full Jacobian by the reduced first-order Jacobian used in the linear multigrid scheme. Using either the non-linear

or the linear multigrid solver as a preconditioner for GMRES, with 20 search directions, produces a sizable increase in speed of convergence, as shown in Figure 5.13. However, the speedup is more pronounced in the case of the linear multigrid solver, where the Krylov method produces a factor of 2.5 increase in overall convergence per cpu time.




FIG. 5.9. *Illustration of Unstructured Grid for Viscous Flow over Three-Element Airfoil (61,104 points)*




FIG. 5.10. *Illustration of Implicit Lines for Viscous Flow over Three-Element Airfoil*




FIG. 5.11. *Computed Mach Contours for Viscous Flow over Three-Element Airfoil*




FIG. 5.12. *Convergence Rates for Linear and Non-Linear Multigrid Algorithms in Terms of Outer Non-Linear Iterations for 3-Element Airfoil Flow Problem*




FIG. 5.13. *Convergence Efficiency for Various Algorithms for 3-Element Airfoil Flow Problem*




FIG. 5.14. *Convergence Efficiency for Various Algorithms for 3-Element Airfoil Flow Problem as Measured by Preconditioned Residual History*

Upon initiating the Krylov method, a large jump in the residuals is observed, which is attributed to the fact that the current Newton-Krylov method operates on the preconditioned residual (i.e. left preconditioning). When the convergence history is plotted in terms of the preconditioned residual, which corresponds to the non-linear corrections produced by the multigrid scheme, a monotone behavior is observed. However, as can be seen from these two plots, conclusions concerning the relative solution efficiency of the various schemes may be a function of the particular measure of convergence.

Table 5.3 illustrates the asymptotic convergence rates achieved by the linear multigrid scheme for all three cases when the linear system is solved to completion at each non-linear update. This represents a lower limit achievable with the multigrid schemes as solvers, and is due to the difference between the true Jacobian of the second order discretization, and the approximate first-order Jacobian employed in the linearization for the linear multigrid scheme, which is also used in the local Jacobians and coarse levels for the FAS multigrid scheme.

This rate is seen to be substantially slower for the last case, indicating that convergence difficulties in this case cannot be addressed through improved linear multigrid methods or, for that matter, any linear

solver based on the first-order Jacobian. Hence, improved restriction, prolongation, coarse grid operators, or agglomeration techniques will have little effect in this case, and future research should concentrate either on better full Jacobian approximations, or improved Krylov methods.

TABLE 5.3
*Asymptotic Convergence Rates Observed for Various
Cases when Linear System Solved to Completion at Each
Non-Linear Iteration*

Case	Asymptotic Rate
NACA 0012 Euler Transonic	0.78
RAE 2822 NS Transonic	0.76
Multi-Element NS Subsonic	0.965

6. Conclusions. The preceding examples demonstrate how linear multigrid methods can deliver superior asymptotic convergence efficiency over non-linear multigrid methods for fluid flow or radiation diffusion problems. When exact Jacobians are available, similar asymptotic convergence rates per multigrid cycle are observed for equivalent linear and non-linear multigrid methods. The efficiency gains of the linear methods are largely attributed to the reduced number of costly non-linear residual evaluations required, and the ability to employ a linear Gauss-Seidel smoother in the place of a Jacobi smoother. Therefore, in cases where costly or complicated non-linear discretizations are employed, the use of linear methods can be advantageous.

Additional convergence acceleration can be achieved by using both linear and non-linear methods as preconditioners to a Newton-Krylov method. This approach is particularly beneficial in cases where an inaccurate linearization is employed by the multigrid solvers.

These conclusions only apply to the solution efficiency in regions of monotonic asymptotic non-linear convergence and when globalization methods are not required. While many practical cases exist (particularly for time-dependent problems) where this behavior is observed, the issues of non-linear convergence and robustness have not been addressed herein, and may affect the performance of a non-linear method over a linear method. Furthermore, the required Jacobian storage for the linear multigrid approach can be prohibitive for many applications, particularly in three dimensions.

This study is to be extended into three dimensions in the near future, and to parallel computer environments. While the overall comparisons can be expected to be similar in three dimensions, evidence shows that linear methods may suffer more efficiency degradation on parallel machines due to the larger number of cheaper grid sweeps employed, which has the effect of raising the computation to communication ratio [24].

REFERENCES

- [1] A. BRANDT, *Multigrid techniques with applications to fluid dynamics:1984 guide*, in VKI Lecture Series, Mar. 1984, pp. 1–176.
- [2] P. N. BROWN AND C. S. WOODWARD, *Preconditioning strategies for fully implicit radiation diffusion with material energy transfer*, Tech. Report UCRL-JC-139087, Lawrence Livermore National Lab, May 2000.
- [3] G. CARRE, *An implicit multigrid method by agglomeration applied to turbulent flows*, Computers and Fluids, 26 (1997), pp. 299–320.

- [4] W. HACKBUSH, *Multigrid Methods and Applications*, Springer-Verlag, Berlin, Germany, 1985.
- [5] A. JAMESON, *Solution of the Euler equations by a multigrid method*, Applied Mathematics and Computation, 13 (1983), pp. 327–356.
- [6] D. A. KNOLL, W. J. RIDER, AND G. L. OLSON, *An efficient nonlinear solution method for nonequilibrium radiation diffusion*, J. Quant. Spec. and Rad. Trans., 63 (1999), pp. 15–29.
- [7] M. LALLEMAND, H. STEVE, AND A. DERVIEUX, *Unstructured multigriding by volume agglomeration: Current status*, Computers and Fluids, 21 (1992), pp. 397–433.
- [8] L. MARTINELLI AND A. JAMESON, *Validation of a multigrid method for the Reynolds-averaged Navier-Stokes Equations*. AIAA Paper 88-0414, Jan. 1988.
- [9] D. J. MAVRIPLIS, *Multigrid techniques for unstructured meshes*, in VKI Lecture Series VKI-LS 1995-02, Mar. 1995.
- [10] ———, *Multigrid strategies for viscous flow solvers on anisotropic unstructured meshes*, Journal of Computational Physics, 145 (1998), pp. 141–165.
- [11] ———, *On convergence acceleration techniques for unstructured meshes*. AIAA Paper 98-2966, presented at the 29th AIAA Fluid Dynamics Conference, Albuquerque, NM, June 1998.
- [12] ———, *Directional agglomeration multigrid techniques for high-Reynolds number viscous flows*, AIAA Journal, 37 (1999), pp. 1222–1230.
- [13] D. J. MAVRIPLIS AND S. PIRZADEH, *Large-scale parallel unstructured mesh computations for 3D high-lift analysis*, AIAA Journal of Aircraft, 36 (1999), pp. 987–998.
- [14] V. A. MOUSSEAU, D. A. KNOLL, AND W. J. RIDER, *Physics-based preconditioning and the Newton-Krylov method for non-equilibrium radiation diffusion*, Journal of Computational Physics, 160 (2000), pp. 743–765.
- [15] E. J. NIELSEN, W. K. ANDERSON, R. W. WALTERS, AND D. E. KEYES, *Application of Newton-Krylov methodology to a three-dimensional unstructured Euler code*, in Proceedings of the 12th AIAA CFD Conference, San Diego, CA, June 1995. AIAA Paper 95-1733-CP.
- [16] N. PIERCE, M. GILES, A. JAMESON, AND L. MARTINELLI, *Accelerating three-dimensional Navier-Stokes calculations*, in Proceedings of the 13th AIAA CFD Conference, Snowmass Village, CO, June 1997. AIAA Paper 97-1953.
- [17] D. J. T. PULLIAM AND P. BUNING, *Recent enhancements to OVERFLOW*. AIAA Paper 97-0644, Jan. 1997.
- [18] M. RAW, *Robustness of coupled algebraic multigrid for the Navier-Stokes equations*. AIAA Paper 96-0297, Jan. 1996.
- [19] P. L. ROE, *Approximate Riemann solvers, parameter vectors and difference schemes*, J. Comp. Phys., 43 (1981), pp. 357–372.
- [20] J. W. RUGE AND K. STÜBEN, *Algebraic multigrid*, in Multigrid Methods, S. F. McCormick, ed., SIAM Frontiers in Applied Mathematics, Philadelphia, 1987, SIAM, pp. 73–131.
- [21] Y. SAAD, *Iterative Methods for Sparse Linear Systems*, PWS Series in Computer Science, PWS Publishing Company, Boston, MA, 1996.
- [22] W. A. SMITH, *Multigrid solution of transonic flow on unstructured grids*, in Recent Advances and Applications in Computational Fluid Dynamics, Nov. 1990. Proceedings of the ASME Winter Annual Meeting, O. Baysal, ed.
- [23] P. R. SPALART AND S. R. ALLMARAS, *A one-equation turbulence model for aerodynamic flows*, La Recherche Aéronautique, 1 (1994), pp. 5–21.

- [24] L. STALS, *The parallel solution of radiation transport equations*, in Proc. of the Tenth SIAM Conference on Parallel Processing for Scientific Computing, Portsmouth, VA, Mar. 2001.
- [25] J. L. THOMAS, D. L. BONHAUS, AND W. K. ANDERSON, *An $(O)(nm^2)$ plane solver for the compressible Navier-Stokes equations*. AIAA Paper 99-0785, 37th AIAA Aerospace Sciences Meeting, Reno NV, Jan. 1999.
- [26] E. TURKEL, *Preconditioning-squared methods for multidimensional aerodynamics*, in Proceedings of the 13th AIAA CFD Conference, Snowmass, CO, June 1997, pp. 856–866. AIAA Paper 97-2025-CP.
- [27] B. VAN LEER, W. T. LEE, AND P. ROE, *Characteristic time-stepping or local preconditioning of the Euler equations*, in Proceedings of the 10th AIAA CFD Conference, Honolulu, HI, June 1991, pp. 260–282. AIAA Paper 91-1552-CP.
- [28] B. VAN LEER, C. H. TAI, AND K. G. POWELL, *Design of optimally-smoothing multi-stage schemes for the Euler equations*. AIAA Paper 89-1933, June 1989.
- [29] V. VENKATAKRISHNAN AND D. J. MAVRIPLIS, *Implicit solvers for unstructured meshes*, Journal of Computational Physics, 105 (1993), pp. 83–91.
- [30] ———, *Agglomeration multigrid for the three-dimensional Euler equations*, AIAA Journal, 33 (1995), pp. 633–640.
- [31] L. B. WIGTON, N. J. YU, AND D. P. YOUNG, *GMRES acceleration of computational fluid dynamic codes*, in Proceedings of the 7th AIAA CFD Conference, July 1985, pp. 67–74. AIAA Paper 85-1494-CP.